Assignment Games with Conflicts: Price of Total Anarchy and Convergence Results via Semi-Smoothness*

Elliot Anshelevich John Postl Tom Wexler

April 19, 2013

Abstract

We study assignment games in which jobs select machines, and in which certain pairs of jobs may conflict, which is to say they may incur an additional cost when they are both assigned to the same machine, beyond that associated with the increase in load. Questions regarding such interactions apply beyond allocating jobs to machines: when people in a social network choose to align themselves with a group or party, they typically do so based upon not only the inherent quality of that group, but also who amongst their friends (or enemies) choose that group as well. We show how *semi-smoothness*, a recently introduced generalization of smoothness, is necessary to find tight or near-tight bounds on the price of total anarchy, and thus on the quality of correlated and Nash equilibria, for several natural job-assignment games with interacting jobs. For most cases, our bounds on the price of total anarchy are either exactly 2 or approach 2. We also prove new convergence results implied by semi-smoothness for our games. Finally we consider coalitional deviations, and prove results about the existence and quality of Strong equilibrium.

1 Introduction

In many real-world scheduling problems, the cost of assigning a given job to a machine depends not only on the total load of that machine, but also on the particular other jobs assigned there. Certain pairs of jobs may conflict, which is to say they may incur an additional cost when they are both assigned to the same machine, beyond that associated with the increase in load. Such conflicts may arise, e.g., when two jobs both place intensive demands on the same computational resources, such as memory, CPU, or network access – assigning one CPU-intensive job and one memory-intensive job to a single machine is often preferable to assigning two jobs that are both memory-intensive, as the former allows for a more uniform utilization of resources. Conversely, some pairs of jobs may benefit from being assigned to the same machine. For example, consider two related processes that need to exchange data during their execution. By assigning these to a single machine, the additional communication cost of sending messages over a network may be avoided.

Load balancing – the efficient allocation of jobs to machines – is one of the most fundamental algorithmic problems in the areas of distributed computing and network design. The associated classes of games, in which players assign their own jobs to machines, are among the most well-studied within the field of algorithmic game theory (see [6, 8, 20, 29, 30] and their references for surveys of this expansive literature). Nevertheless, interactions between jobs like the ones described

*This work was supported in part by NSF grants CCF-0914782, CCF-1101495, and CNS-1017932.

above have not been extensively analyzed, even for very simple contexts. Questions regarding such interactions apply beyond allocating jobs to machines: when people in a social network choose to align themselves with a group or party, they typically do so based upon not only the inherent quality of that group, but also on who amongst their friends (or enemies) chooses that group as well. In this setting groups correspond to machines and people correspond to jobs. Although we will usually talk about jobs and machines in this paper, the questions we ask can be framed in the context of social networks as well.

In this paper we aim to model and analyze the effects of such interactions. To this end, we consider two basic job-assignment games, and superimpose a graph structure on jobs that represents their pairwise interactions. In the resulting games, a given player's utility depends both on total load of their selected machine and the placement of jobs by their neighbors in this interaction graph. Despite the simplicity of the underlying models, the addition of this structure yields rich classes of games, fundamentally distinct from the originals, with new properties and non-trivial behavior. Our primary goal is to understand and quantify the efficiency of these games.

Perhaps the most widely studied characterization of game efficiency is the price of anarchy [21] – the ratio of the social cost of the worst Nash equilibrium to that of an optimal, centrally designed solution. As the field of algorithmic game theory has matured, more attention has been paid to the evaluation of additional equilibrium concepts. Some of the most notable of these concepts are correlated and coarse-correlated equilibrium, the latter being especially important due to its connection with outcomes of no-regret learning [27]. The ratio between the quality of the worst coarse-correlated equilibrium and the centrally optimal solution is sometimes called the *price of total anarchy* [5, 27] (this is how we will use it in this paper, although this concept was originally defined as a bound on the results of no-regret learning). The price of total anarchy trivially acts as a bound on the quality of correlated and Nash equilibrium as well. In the unifying paper of Roughgarden [27], the concept of smoothness is introduced, which provides a general framework for bounding the price of total anarchy, as well as for generating convergence results for best-response dynamics. However, as we show below, the standard version of this smoothness framework cannot yield tight bounds on the quality of equilibria for our games.

We instead show how *semi-smoothness*, a generalization of smoothness recently introduced in [22] for the analysis of a GSP game, is necessary to form such tight or near-tight bounds. In particular, semi-smoothness with respect to a mixed strategy is needed to show a tight bound in our games. We also prove new convergence results implied by semi-smoothness for our games. Our results illustrate the usefulness of semi-smoothness as a proof technique for natural games beyond GSP.

1.1 Our Contributions

We use semi-smoothness to bound the quality of coarse-correlated equilibrium (CCE) for several natural job-assignment games with interacting jobs. We show that these bounds are tight (or near-tight), not only for CCE but for correlated and Nash equilibria as well. We also prove new convergence results for best-response dynamics in these games. Specifically, we consider the following games (as well as their extensions: see Sections 3.3 and 4.2).

In the game that we call *Balancing with Conflicts* (BwC), the strategy space of each of the n players consists of $\{1, \ldots, m\}$. We can think of each player as a job choosing among m machines, or as a person choosing among m groups. In the usual (and thoroughly analyzed) load balancing game, the cost to a player choosing machine k is an increasing function of $|X_k|$, where X_k is the set of players who choose machine k. In the BwC game we are also given an undirected graph $G = (V, E)$ in which nodes correspond to players and an edge represents a conflict between these

players. Two players who are neighbors in G and are assigned to the same machine both incur an additional cost. Specifically, the cost of a player i who selects strategy k is

$$c_i = |X_k| + |X_k \cap \{j|(i,j) \in E\}|.$$

We define the cost of a solution as the sum of player costs, i.e., $c(\mathbf{s}) = \sum_i c_i(\mathbf{s})$. Even in this simple setting, the addition of conflict costs drastically changes the game. Using semi-smoothness, we show that all CCE are within a factor of $2 - \frac{1}{m} + \frac{m-1}{n}$ ($\frac{3}{2} + \frac{1}{n}$ for 2 machines) of the optimum solution (we suspect the actual bound is $2 - \frac{1}{m}$), as well as provide new convergence results. Note that this bound cannot be obtained using arguments based on smoothness. In addition to various extensions of this game, we prove that strong Nash equilibrium, which is a solution that is stable with respect to coalitional deviations [26], always exists for BwC, and that the strong price of anarchy is at most $\frac{4}{3} + \frac{2}{3n}$ for $m = 2$. Note that games in which people partition themselves into two groups while trying to avoid being with others they dislike is already a nontrivial and important special case [4,14]. All of the above results also hold for the version of this game where players attempt to minimize the number of neighbors *not* in their group (the player cost is $|X_k| + |\{j|(i,j) \in E\} \backslash X_k|$), i.e., players prefer to group with as many "friends" as possible.

We also study a utility-maximizing game which we call *Sharing with Conflicts* (SwC). In this game the strategy space of each of the n players still consists of $\{1, \ldots, m\}$, but each machine/group k has an intrinsic value (or processing power) p_k. This value is equally shared among the players who choose machine k. Without conflicts, this is a variation on the market sharing game of, e.g., [2,13,23], but once again we suppose that players benefit from avoiding those machines selected by their neighbors in some network G. Since this is a utility-maximization game, we make this explicit by setting the utility of a player i who selects strategy k to be

$$u_i = p_k/|X_k| + |\{j|(i,j) \in E\} \backslash X_k|.$$

The social welfare $u(\mathbf{s})$ is defined as the sum of player utilities, i.e., $u(\mathbf{s}) = \sum_i u_i(\mathbf{s})$. We use the semi-smoothness framework to show new convergence results for this and related games, as well as a tight bound of 2 for the quality of coarse correlated equilibrium as compared to the optimum solution. These results also hold for weighted edges, where the utility of being assigned to different machines can be different for each pair of players.

1.2 Related Work

Our games are related to many well-studied games. There is a significant amount of literature on the subject of selfish load balancing and, more generally, atomic congestion games; consult, for example, [6, 8, 20, 29, 30] and their references. In the market sharing game [2, 13, 23], there are locations with values that are divided equally among all players at that particular location. Cut games, such as max cut, are represented by graphs in which the nodes correspond to players and the edges signify that two players are either friends or enemies. Although our games are very related (see Section 5), our results cannot be obtained from known results about cut games. Hoefer [16] gives price of anarchy bounds for the max k-cut game. Gourvès and Monnot [14, 15] explored the existence and inefficiency of strong Nash equilibria for the max k-cut game. We use techniques similar to [15] to prove strong price of anarchy bounds for our games. Since max cut is PLS-complete, efforts have been focused on finding good solutions. For example, Bhalgat et al. [4] give an algorithm for finding approximate Nash equilibrium and Christodoulou et al. [9] show that random walks can converge to good approximate solutions very quickly.

There are many games that model the idea of players wanting to be close to their friends and away from their enemies, which is closely related to the motivation behind our games. The

aforementioned cut games model this concept. Clustering games, such as those found in [12, 16], are clearly representative of this idea as well. Coordination games in which players choose whether or not to adopt a new technology depending on what their friends are doing [17–19, 25] are also related to this concept, since players are essentially choosing to be in the same group with as many of their friends as possible. However, despite the similarity in motivation, the models and issues arising in this line of work are very different from ours.

The concept of smoothness was originally defined by Roughgarden [27]. Smoothness can be used to bound the price of anarchy for many equilibrium concepts, and it implies fast convergence to good solutions. Several variants of smoothness have been introduced [7,22,28] to address the inability of the original smoothness definition to provide tight price of anarchy bounds for various games. We use the variant in [7], known as *semi-smoothness*, for this reason. The concept of β-niceness was originally defined by Awerbuch, et al. [2] and explored again in [1]. Like smoothness, β-niceness provides bounds for price of anarchy, but only for pure Nash equilibria. Furthermore, a β-nice game that has a potential function (see [26] for the definition of "potential games") and satisfies some other conditions will converge to states that are approximately optimal (although not necessarily stable) in a polynomial number of steps [1,2]. Other papers, such as [10,11,27], show that potential games can converge to good states quickly under certain conditions. We demonstrate that a more general variant of β-niceness, called (λ, μ)-*niceness*, is simply a weaker version of smoothness. Our convergence results generalize those from [1, 27]; we only require that they be potential games and (λ, μ)-nice. The games do not need to satisfy a bounded-jump condition like in [1, 2], or be smooth like in [27].

2 Smoothness Concepts and Preliminaries

In this section we present how semi-smoothness relates to various similar concepts, and establish some of its implications. We begin by defining these concepts for payoff-maximization games; the definitions for cost-minimization games are similar, and are included in Section 2.3 for completeness. Throughout this paper, $\mathbf{s}^*$ will refer to a centrally optimal solution.

2.1 Payoff-Maximization Games

First recall the classic definition of smoothness from [27]. Smoothness of a game implies many positive results about it, including price of anarchy bounds for various equilibrium concepts and good convergence results.

Definition 2.1. (Smooth Games) A payoff-maximization game is (λ, μ)-*smooth* if for every pair of outcomes $\mathbf{s}$ and $\mathbf{s}'$,

$$\sum_{i\in N} u_i(s_i', s_{-i}) \geq \lambda \cdot u(\mathbf{s}') - \mu \cdot u(\mathbf{s}).$$

In the above definition, $u_i(\mathbf{s})$ is the utility of player i in outcome $\mathbf{s}$, and $u(\mathbf{s}) = \sum_i u_i(\mathbf{s})$. Most relevant to our work is the fact that a game being (λ, μ)-smooth with nonnegative λ, μ implies that the total utility of all coarse correlated equilibria (CCE) is at least $\frac{\lambda}{1+\mu}$ times the utility of the optimum solution (the outcome with the highest total utility). This immediately implies a similar bound for outcomes that result from no-regret learning, and for both pure and mixed Nash equilibrium [27]. Thus, the price of total anarchy [5] is at most $\frac{1+\mu}{\lambda}$.

Other, more general notions of smoothness have also been studied [7, 22, 28]. Here we mention the one most relevant to our work, known as semi-smoothness.

Definition 2.2. (Semi-Smooth Games) A payoff-maximization game is (λ, μ)-*semi-smooth* if there exists a mixed strategy σ_i for each player i such that for every outcome $\mathbf{s}$,

$$\mathbf{E}_\sigma \left[\sum_{i \in N} u_i(\sigma_i, s_{-i}) \right] \geq \lambda \cdot u(\mathbf{s}^*) - \mu \cdot u(\mathbf{s}).$$

It is easy to show that a game being semi-smooth implies all the same results from [27] as smoothness, including a bound of $\frac{1+\mu}{\lambda}$ for the price of total anarchy. Notice that it essentially removes one of the '$\forall$' quantifiers: while smoothness requires that the inequality holds for all $\mathbf{s}, \mathbf{s}'$, semi-smoothness only requires that there *exist* a mixed outcome σ such that for all $\mathbf{s}$, the inequality holds. Thus, semi-smoothness is quite a bit more general.

Semi-smoothness was first defined in [22], although in [22] σ was restricted to pure strategies. It was later re-defined with mixed strategies in [7]. The strategies σ_i being mixed makes a large difference when establishing bounds on the price of total anarchy. While in the later sections of this paper we study games where this provides tight bounds on the price of total anarchy, here we give a simple example of a well-known game where semi-smoothness with mixed strategies provides a tight bound, while the standard definition of smoothness (or even semi-smoothness with σ_i being pure strategies) does not. Specifically, we consider the well-studied Max-Cut game:

Definition 2.3. (Max-Cut Game) We are given a graph $G = (V, E)$; each player corresponds to a node in G. Each player must choose to be in partition 1 or partition 2. The utility of player i is the number of edges between i and players that choose a different partition from i.

It is well-known that the price of anarchy of the Max-Cut Game is 2. However, as we show below, there is no choice of λ and μ for which this game is (λ, μ)-smooth (or semi-smooth with pure strategies) for any λ, μ such that $\frac{\lambda}{1+\mu} > \frac{1}{3}$. On the other hand, the Max-Cut game is $(\frac{1}{2}, 0)$-semi-smooth, thus implying not only that the price of anarchy is 2, but also that the utility of any CCE is within a factor of 2 of the optimum.

Theorem 2.4. The Max-Cut Game is $(\frac{1}{2}, 0)$-semi-smooth. On the other hand, there is no choice of λ and μ for which this game is (λ, μ)-smooth (or semi-smooth with pure strategies) such that $\frac{\lambda}{1+\mu} > \frac{1}{3}$.

Proof. Let a graph $G = (V, E)$ be an instance of the max-cut game, and let σ_i be the mixed strategy where player i chooses each of the two partitions with probability $\frac{1}{2}$. For all outcomes $\mathbf{s}$, $u_i(1, s_{-i}) + u_i(2, s_{-i})$ equals the total degree of node i, which we denote by d_i. Thus, we observe that

$$\mathbf{E} \left[\sum_{i \in N} u_i(\sigma_i, s_{-i}) \right] = \frac{1}{2} \sum_{i \in N} d_i = |E| \geq \frac{1}{2} u(\mathbf{s}^*).$$

Thus, the max-cut game is $\left(\frac{1}{2}, 0\right)$-*semi-smooth.*

Now we show that the same bound cannot be proven using previously studied concepts of smoothness. Consider a graph G which contains two nodes i, j and a single edge between them. Let $\mathbf{s}^1$ be an optimal solution, i.e., $s_i^1 = 1, s_j^1 = 2$. Then $u(\mathbf{s}^1) = 2$. Let $\mathbf{s}^2$ be a permutation of $\mathbf{s}^1$, i.e., $s_i^2 = 2, s_j^2 = 1$; $u(\mathbf{s}^2) = 2$. Let $\mathbf{s}^3$ be the solution in which both players are in partition 1 and $\mathbf{s}^4$ be the solution in which both players are in partition 2. Then $u(\mathbf{s}^3) = u(\mathbf{s}^4) = 0$. We must show that for any choice of σ such that σ is a vector of pure strategies, max cut is not (λ, μ)-semi-smooth with pure strategies for any λ, μ such that $\frac{\lambda}{1+\mu} > \frac{1}{3}$.

Let $\sigma = \mathbf{s}^1$. Consider when players deviate from $\mathbf{s}^2$. If player i deviates from s_i^2 to s_i^1 without player j changing his strategy from s_j^2, then player i's utility becomes 0. Similarly, if player j deviates from s_j^2 to s_j^1, then player j's utility becomes 0. That is, $u_i(\sigma_i, s_{-i}^2) + u_j(\sigma_j, s_{-j}^2) = 0$. We observe that $0 \geq \lambda \cdot u(\mathbf{s}^1) - \mu \cdot u(\mathbf{s}^1) = 2(\lambda - \mu)$ iff $\lambda \leq \mu$. Next, we consider when players deviate from $\mathbf{s}^3$. Player i's utility remains 0 if she deviates to $\mathbf{s}_i^1$ from $\mathbf{s}_i^3$ since $\mathbf{s}_i^1 = \mathbf{s}_i^3$. However, if player j deviates to $\mathbf{s}_j^1$ from $\mathbf{s}_j^3$, then her utility becomes 1. That is, $u_i(\sigma_i, s_{-i}^3) + u_j(\sigma_j, s_{-j}^3) = 1$. We observe that $1 \geq \lambda \cdot u(\mathbf{s}^1) - \mu \cdot u(\mathbf{s}^3)$ iff $\lambda \leq \frac{1}{2}$. Thus, when $\sigma = \mathbf{s}^1$, $\frac{\lambda}{1+\mu} \leq \frac{1}{3}$, since $\lambda \geq 0$, $\lambda \leq \mu$ and $\lambda \leq \frac{1}{2}$. Similar analysis holds for $\sigma = \mathbf{s}^2$.

Let $\sigma = \mathbf{s}^3$. Consider when players deviate from $\mathbf{s}^3$. Since the players are deviating to the same strategy in which their utilities were 0, their utilities remain 0. That is, $u_i(\sigma_i, s_{-i}^3) + u_j(\sigma_j, s_{-j}^3) = 0$. We observe that $0 \not\geq \lambda \cdot u(\mathbf{s}^1) - \mu \cdot u(\mathbf{s}^3)$ for any choice of nonnegative λ, μ. Similar analysis holds for $\sigma = \mathbf{s}^4$.

We have considered every possible σ. We conclude that max cut is not (λ, μ)-semi-smooth with pure strategies for any λ, μ such that $\frac{\lambda}{1+\mu} > \frac{1}{3}$. From this, we can also conclude that max cut is not (λ, μ)-smooth for any λ, μ such that $\frac{\lambda}{1+\mu} > \frac{1}{3}$. ■

Finally, we introduce another related notion of smoothness. This last definition of smoothness is actually a modified version of *β-niceness* from [1] and [2]. It is much weaker than the previous versions; while semi-smoothness still implies all the same results as the usual concept of smoothness (including convergence results, as we discuss below), niceness only implies bounds for Nash equilibrium and price of anarchy, and says nothing about CCE or price of total anarchy. However, it is sufficient to provide good convergence results.

Definition 2.5. (Nice Games) A payoff-maximization game is (λ, μ)-*nice* if for every outcome $\mathbf{s}$, there exists an outcome $\mathbf{s}'$ such that

$$\sum_{i \in N} u_i(s_i', s_{-i}) \geq \lambda \cdot u(\mathbf{s}^*) - \mu \cdot u(\mathbf{s}).$$

The difference between nice and semi-smooth games is that $\mathbf{s}'$ is allowed to depend on $\mathbf{s}$, while in semi-smooth games σ must be a single mixed state, independent of $\mathbf{s}$. It is easy to verify that while niceness implies a bound of $\frac{1+\mu}{\lambda}$ on price of anarchy, to show the same results for CCE or no-regret outcomes, the solution that the players are "switching to" must be independent of $\mathbf{s}$. It follows from their definitions that for fixed λ, μ, (λ, μ)-*smooth games* $\subseteq$ (λ, μ)-*semi-smooth games* $\subseteq$ (λ, μ)-*nice games*.

2.2 Convergence Results for Payoff-Maximization Games

Now that we have defined various concepts related to smoothness, and established which of these concepts imply bounds on the quality of CCE and/or Nash equilibrium, we proceed to describe convergence results, which only require that a game is *nice*, and therefore also hold for all the concepts mentioned above. Note that as in [1, 27], these results only hold for potential games [24]. We define $\Delta_i(\mathbf{s}) = u_i(s_i^b, s_{-i}) - u_i(\mathbf{s})$, where s_i^b is the best response of player i to state s, as the amount of utility player i gains from deviating, and $\Delta(\mathbf{s}) = \sum_{i \in N} \Delta_i(\mathbf{s})$. When we refer to the BR dynamics, we will mean a sequence of best response moves in which the player selected to make their move is the one with the maximum improvement $\Delta_i(\mathbf{s})$ [1]. Similar results can be shown for many other types of best response dynamics.

The first easy convergence result is a slight generalization of the one given by [27], with the proof being essentially the same. It works for arbitrary potential functions instead of being restricted to lower potential functions, as in [27]. As usual, $\mathbf{s}^*$ refers to the solution with highest utility.

Theorem 2.6. Consider any payoff-maximization game H that has an exact potential function $\Phi(\mathbf{s})$ such that for some $B \geq 1$, we have that $u(\mathbf{s}) \geq \frac{1}{B}\Phi(\mathbf{s})$ and is (λ, μ)-*nice*. Let $\rho = \frac{\lambda}{1+\mu}$. Then, for any $\epsilon > 0$, the BR dynamic reaches a state $\mathbf{s}^t$ with $u(\mathbf{s}^t) \geq \frac{\rho}{1+\epsilon} \cdot u(\mathbf{s}^*)$ from any starting state $\mathbf{s}^0$ in at most $O\left(\frac{Bn}{\epsilon(1+\mu)} \log(B \cdot u(\mathbf{s}^*))\right)$ steps.

Note that the above theorem does *not* say that all states after state s^t will be good, just that a good state will be reached in a small amount of time. Our second convergence result is a generalization of the convergence result for perfect β-*nice* games in [1]. Our version of this result does not require the game to be perfect and allows more general potential functions.

Theorem 2.7. Consider any payoff-maximization game H that has an exact potential function $\Phi(\mathbf{s})$ such that for some $A, B \geq 1$, we have that $A\Phi(\mathbf{s}) \geq u(\mathbf{s}) \geq \frac{1}{B}\Phi(\mathbf{s})$ and is (λ, μ)-*nice*. Let $\rho = \frac{\lambda}{1+\mu}$. Then, for any $\epsilon > 0$, the BR dynamic reaches a state $\mathbf{s}^t$ with $u(\mathbf{s}^t) \geq \frac{\rho(1-\epsilon)}{AB} \cdot u(\mathbf{s}^*)$ in at most $O\left(\frac{n}{A(1+\mu)} \log \frac{1}{\epsilon}\right)$ steps from any initial state. Furthermore, all future states reached with best-response dynamics will satisfy this approximation factor as well.

Proof. Let i be the player selected by the BR dynamic at any given time step. By our hypotheses about $\Phi(\mathbf{s})$ and the definition of (λ, μ)-*nice*,

$$\begin{aligned}
\Phi(s_i^b, s_{-i}) - \Phi(\mathbf{s}) &= u_i(s_i^b, s_{-i}) - u_i(\mathbf{s}) \\
&\geq \frac{\Delta(\mathbf{s})}{n} \\
&\geq \frac{\lambda \cdot u(\mathbf{s}^*) - (1+\mu) \cdot u(\mathbf{s})}{n} \\
&\geq \frac{\lambda \cdot u(\mathbf{s}^*) - A(1+\mu)\Phi(\mathbf{s})}{n}.
\end{aligned}$$

Let $f(\mathbf{s}) = \frac{\lambda \cdot u(\mathbf{s}^*) - A(1+\mu)\Phi(\mathbf{s})}{n}$. Then $f(\mathbf{s}) - f(s_i^b, s_{-i}) = \frac{A(1+\mu)}{n}\left(\Phi(s_i^b, s_{-i}) - \Phi(\mathbf{s})\right) \geq \frac{A(1+\mu)}{n} f(\mathbf{s})$, which implies that $f(s_i^b, s_{-i}) \leq \left(1 - \frac{A(1+\mu)}{n}\right) f(\mathbf{s})$. Thus, from a starting state $\mathbf{s}^0$, the BR dynamic converges to a state $\mathbf{s}^t$ with

$$f(\mathbf{s}^t) \leq \left(1 - \frac{A(1+\mu)}{n}\right)^t f(\mathbf{s}^0).$$

We can set $t = \lceil \frac{n}{A(1+\mu)} \ln \frac{1}{\epsilon} \rceil$. Using the fact that $\left(1 - \frac{1}{x}\right)^x \leq 1/e$, we are able to derive that $f(\mathbf{s}^t) \leq e^{\ln \epsilon^{-1}} f(\mathbf{s}^0) \leq \epsilon \cdot f(\mathbf{s}^0) \leq \frac{\epsilon \lambda \cdot u(\mathbf{s}^*)}{n}$. Finally, using these results and our $\Phi(\mathbf{s})$ bounds, we can derive

$$\begin{aligned}
u(\mathbf{s}^t) &\geq \frac{1}{B}\Phi(\mathbf{s}^t) \\
&= \frac{n}{AB(1+\mu)}\left(\frac{\lambda \cdot u(\mathbf{s}^*)}{n} - f(\mathbf{s}^t)\right) \\
&\geq \frac{n}{AB(1+\mu)}\left((1-\epsilon)\frac{\lambda \cdot u(\mathbf{s}^*)}{n}\right) \\
&\geq \frac{\rho(1-\epsilon)}{AB} \cdot u(\mathbf{s}^*).
\end{aligned}$$

Because $\Phi(\mathbf{s}^t) \geq \frac{\rho(1-\epsilon)}{A} \cdot u(\mathbf{s}^*)$, we know that this approximation factor will hold for all future states reached by best-response dynamics as well, since Φ will only increase. ∎

2.3 Cost-Minimization Games

While above we defined everything for utility-maximization games, the same concepts can be defined and results hold for cost-minimization games. For completeness, we define these concepts here.

Definition 2.8. (Smooth Games) A cost-minimization game is (λ, μ)-*smooth* if for every pair of outcomes $\mathbf{s}$ and $\mathbf{s}'$,

$$\sum_{i \in N} c_i(s_i', s_{-i}) \leq \lambda \cdot c(\mathbf{s}') + \mu \cdot c(\mathbf{s}).$$

Definition 2.9. (Semi-Smooth Games) A cost-minimization game is (λ, μ)-*semi-smooth* if there exists a mixed strategy σ_i for each player i such that for every outcome $\mathbf{s}$,

$$\mathbf{E}_\sigma\left[\sum_{i \in N} c_i(\sigma_i, s_{-i})\right] \leq \lambda \cdot c(\mathbf{s}^*) + \mu \cdot c(\mathbf{s}).$$

If a cost-minimization game is (λ, μ)-*semi-smooth*, then the price of anarchy for coarse correlated equilibria is at most $\frac{\lambda}{1-\mu}$.

Definition 2.10. (Nice Games) A cost-minimization game is (λ, μ)-*nice* if for every outcome $\mathbf{s}$, there exists an outcome $\mathbf{s}'$ such that

$$\sum_{i \in N} c_i(s_i', s_{-i}) \leq \lambda \cdot c(\mathbf{s}^*) + \mu \cdot c(\mathbf{s}).$$

3 Balancing with Conflicts

Definition. Balancing with Conflicts can be described as a triple (N, M, G), where $N = \{1, 2, \ldots, n\}$ is a set of players, $M = \{1, 2, \ldots, m\}$ is a set of machines, and $G = (N, E)$ is an undirected graph on the players. Each player i must select one machine. That is, the strategy set for each player i is $S_i = M$. Intuitively, this models a job being assigned to a machine to be processed. A state of our game is the assignment of every player to a machine, which can be represented by a vector $\mathbf{s} = (s_1, s_2, \ldots, s_n)$ where $s_i \in S_i$ for each player $i \in N$.

Given a state $\mathbf{s}$, let $X_k(\mathbf{s})$ denote the set of players assigned to machine k, and let $x_k(\mathbf{s}) = |X_k(\mathbf{s})|$. Let $E(X, Y) = \{(i, j) \in E : i \in X, j \in Y\}$ denote the edges between the players in sets X and Y, and let $e(X, Y) = |E(X, Y)|$. We define $E(X) = E(X, X)$ with $e(X) = |E(X)|$.

We will now define the cost incurred by a player. For each player, including itself, that is assigned to to the same machine as player i, i incurs a cost of 1. Furthermore, if two jobs on the same machine conflict there is an additional cost, i.e., for each player j such that j shares an edge with player i and j is in the same machine as i, i incurs an additional cost of 1. Formally, for a player i with $s_i = k$, its cost is $c_i(\mathbf{s}) = x_k(\mathbf{s}) + e(\{i\}, X_k(\mathbf{s}))$. Finally, we define the cost of an outcome of our game as the total cost experienced by the players, that is, $c(\mathbf{s}) = \sum_{i \in N} c_i(\mathbf{s}) = \sum_{k=1}^{m} \left(x_k(\mathbf{s})^2 + 2e(X_k(\mathbf{s}))\right)$.

3.1 Near Tight Bounds for Quality of CCE

We will show that Balancing with Conflicts has several desirable properties. First, it is easy to see that Balancing with Conflicts is an exact potential game with $\Phi(\mathbf{s}) = \frac{1}{2}c(\mathbf{s})$. This guarantees that every instance of Balancing with Conflicts admits a pure Nash equilibrium. The existence of a potential function guarantees that best-response dynamics will converge to a stable solution, but this may require exponential time. We will show later that other properties guarantee that best-response dynamics converge to a good (but not necessarily stable) state very quickly.

Since $c(\mathbf{s}^*)$ is the minimum value of the objective function by definition, it follows that $\Phi(\mathbf{s}^*)$ is the global minimum of the potential function. This immediately implies that the optimal solution is a pure Nash equilibrium. Thus, we conclude that the price of stability of the Balancing with Conflicts game is 1.

We will now give the main result of this section, but we will prove it later. Previous smoothness definitions were too strong to find good price of anarchy bounds for Balancing with Conflicts. However, using *semi-smoothness*, we are able to derive asymptotically tight price of anarchy bounds. Furthermore, this result has implications beyond price of anarchy; most notably, it guarantees fast convergence to "good" solutions.

Theorem 3.1. The Balancing with Conflicts game is $(2 - \frac{1}{m} + \frac{m-1}{n}, 0)$*-semi-smooth.*

The following price of anarchy results for pure Nash equilibrium, mixed Nash equilibrium, correlated equilibrium, and coarse correlated equilibrium follow immediately from Theorem 3.1.

Corollary 3.2. The price of total anarchy of the Balancing with Conflicts game is at most $2 - \frac{1}{m} + \frac{m-1}{n}$.

We will now consider how to find good solutions. First, we can trivially create a 2-approximately stable solution by randomly assigning players to machines such that at most $\lceil \frac{n}{m} \rceil$ players are assigned to each machine. Secondly, we can use best-response dynamics from any initial state to reach a good state. In [1], Augustine et al. show that perfect β-nice payoff-maximization games reach approximately optimal states very quickly. One can easily derive a similar theorem for cost games that are perfect (λ, μ)-nice. Since Balancing with Conflicts is $(2 - \frac{1}{m} + \frac{m-1}{n}, 0)$-semi-smooth, it must be $(2 - \frac{1}{m} + \frac{m-1}{n}, 0)$-nice as well. A game is *perfect* if for any solution $\mathbf{s}$ and any improving move s_i' of player i,

$$c(\mathbf{s}) - c(s_i', s_{-i}) \geq \Phi(\mathbf{s}) - \Phi(s_i', s_{-i}) \geq c_i(\mathbf{s}) - c_i(s_i', s_{-i})$$

Since $\Phi(\mathbf{s}) = \frac{1}{2}c(\mathbf{s})$ is an exact potential function, it follows that for every job i and strategy s_i', $c(\mathbf{s}) - c(s_i', s_{-i}) = 2\left(\Phi(\mathbf{s}) - \Phi(s_i', s_{-i})\right) \geq \Phi(\mathbf{s}) - \Phi(s_i', s_{-i}) = c_i(\mathbf{s}) - c_i(s_i', s_{-i})$. Since Balancing with Conflicts is perfect, the following corollary follows from Theorem 1 of [1].

Corollary 3.3. For any instance of the Balancing with Conflicts game and for any $\epsilon > 0$, the BR dynamic reaches a state $\mathbf{s}^t$ with

$$c(\mathbf{s}^t) \leq \left(2 - \frac{1}{m} + \frac{m-1}{n}\right)(1+\epsilon) \cdot c(\mathbf{s}^*)$$

in at most $O\left(n \log \frac{m}{\epsilon}\right)$ steps from any initial state $\mathbf{s}^0$. This approximation factor holds for all further states reached by best-response dynamics.

We will now focus on proving Theorem 3.1. We need lower bounds for the optimal solution in order to do so. This first lower bound for all solutions follows from the fact that the load balancing cost is minimized when $\frac{n}{m}$ jobs are assigned to each machine.

Lemma 3.4. For any solution $\mathbf{s}$, $c(\mathbf{s}) \geq \frac{n^2}{m}$.

Proof. Let $\mathbf{x} = (x_1(\mathbf{s}), x_2(\mathbf{s}), \ldots, x_m((\mathbf{s})))$. Let $\|\mathbf{x}\|_1, \|\mathbf{x}\|_2$ denote the L_1 and L_2 norms, respectively. It is known that for any m dimensional vector $\mathbf{x}$, $\|\mathbf{x}\|_1 \leq \sqrt{m} \cdot \|\mathbf{x}\|_2$. Using this property, we can derive

$$\sum_{k=1}^{m} x_k(\mathbf{s})^2 = \|\mathbf{x}\|_2^2 \geq \left(\frac{\|\mathbf{x}\|_1}{\sqrt{m}}\right)^2 = \frac{\left(\sum_{k=1}^{m} x_k(\mathbf{s})\right)^2}{m} = \frac{n^2}{m}.$$

Thus, $c(\mathbf{s}) = \sum_{k=1}^{m} \left(x_k(\mathbf{s})^2 + 2e(X_k(\mathbf{s}))\right) \geq \frac{n^2}{m}$. ■

This next lower bound makes use of the fact that if G has a large enough number of edges, then the solution will unavoidably incur some edge cost.

Lemma 3.5. For any solution $\mathbf{s}$, $(m-1)\cdot c(\mathbf{s}) \geq 2|E|$.

Proof. We begin by considering, for any pair of machines k, l such that $k \neq l$, the number of edges between the set of players on machine k and the set of players on machine l. The number of edges between these sets is maximized when every player on machine k shares an edge with every player on machine l. That is, $e(X_k(\mathbf{s}), X_l(\mathbf{s})) \leq x_k(\mathbf{s}) \cdot x_l(\mathbf{s})$. Using this fact and Lemma 3.4,

$$\begin{aligned}
\sum_{k=1}^{m-1}\sum_{l=k+1}^{m} e(X_k(\mathbf{s}), X_l(\mathbf{s})) &\leq \sum_{k=1}^{m-1}\sum_{l=k+1}^{m} x_k(\mathbf{s})\cdot x_l(\mathbf{s}) \\
&= \frac{1}{2}\left(\sum_{k=1}^{m}\sum_{l=1}^{m} x_k(\mathbf{s})\cdot x_l(\mathbf{s}) - \sum_{k=1}^{m} x_k(\mathbf{s})^2\right) \\
&= \frac{1}{2}\left(n^2 - \sum_{k=1}^{m} x_k(\mathbf{s})^2\right) \\
&\leq \frac{1}{2}\left(n^2 - \frac{n^2}{m}\right) \\
&= \frac{(m-1)\cdot n^2}{2m}.
\end{aligned}$$

The number of edges between players on different machines plus the number of edges between players on the same machines is equal to the total number of edges in the graph G. That is, $|E| = \sum_{k=1}^{m-1}\sum_{l=k+1}^{m} e(X_k(\mathbf{s}), X_l(\mathbf{s})) + \sum_{k=1}^{m} e(X_k(\mathbf{s}))$. Using this fact, our previously derived result, and Lemma 3.4,

$$\begin{aligned}
c(\mathbf{s}) &= \sum_{k=1}^{m}\left(x_k(\mathbf{s})^2 + 2e(X_k(\mathbf{s}))\right) \\
&\geq \frac{n^2}{m} + 2\left(|E| - \sum_{k=1}^{m-1}\sum_{l=k+1}^{m} e(X_k(\mathbf{s}), X_l(\mathbf{s}))\right) \\
&\geq 2|E| + \frac{(2-m)\cdot n^2}{m} \\
&\geq 2|E| + (2-m)\cdot c(\mathbf{s}).
\end{aligned}$$

Rearranging the terms completes the proof. ∎

We finally have the necessary lower bounds, and we will now proceed with the proof of our primary result.

Proof of Theorem 3.1. Let $\mathbf{s}$ be an outcome, and let d_i denote the degree of player i in G. For each player i with $s_i = k$, $c_i(\mathbf{s}) = x_k(\mathbf{s}) + e(\{i\}, X_k(\mathbf{s}))$. If player i deviates to a different machine l, then their cost will be the sum of number of players already on machine l plus 1 for itself and the number of edges from i to players on machine l. That is, for each $s_i' = l \neq s_i$, $c_i(s_i', s_{-i}) = x_l(\mathbf{s}) + 1 + e(\{i\}, X_l(\mathbf{s}))$. For each player i, let σ_i denote the mixed strategy in which every strategy $s_i' \in S_i$ is selected with probability $\frac{1}{m}$. Then for every player i,

$$\mathbf{E}\left[c_i(\sigma_i, s_{-i})\right] = \frac{1}{m}\left(\sum_{k=1}^{m}\left(x_k(\mathbf{s}) + e(\{i\}, X_k(\mathbf{s}))\right) + m - 1\right) = \frac{1}{m}\left(n + d_i + m - 1\right).$$

By linearity of expectation and our previous equation,

$$\begin{aligned}\mathbf{E}\left[\sum_{i\in N} c_i(\sigma_i, s_{-i})\right] &= \sum_{i\in N} \mathbf{E}\left[c_i(\sigma_i, s_{-i})\right] \\ &= \frac{1}{m}\sum_{i\in N}(n + d_i + m - 1) \\ &= \frac{1}{m}\left(n^2 + 2|E| + (m-1)\cdot n\right) \\ &= \left(1 + \frac{m-1}{n}\right)\frac{n^2}{m} + \frac{2|E|}{m}\end{aligned}$$

Finally, we can use our lower bounds for OPT from Lemmas 3.4 and 3.5 to show that this quantity is at most

$$\left(2 - \frac{1}{m} + \frac{m-1}{n}\right)\cdot c(\mathbf{s}^*).$$

■

We now give a lower bound on the price of anarchy that tends to the upper bound as the number of players grows.

Claim 3.6. The price of anarchy of Balancing with Conflicts is at least $2 - \frac{1}{m}$.

Proof. Consider the case where there are m machines and G is a complete m-partite graph where each of the m disjoint sets has m nodes. Then $n = m^2$.

Consider the solution $\mathbf{s}$ formed by assigning the jobs in the first disjoint set to the first machine, the jobs in the second disjoint set to the second machine, etc. Then for each machine k, $x_k(\mathbf{s}) = m$. Then the cost of this solution is $m^3 = \frac{n^2}{m}$. Since $c(\mathbf{s}^*) \geq \frac{n^2}{m}$, $\mathbf{s}$ must be the optimal solution.

Now consider the solution $\mathbf{s}'$ formed by doing the following: for each disjoint set, assign all of the jobs in that set to different machines. We claim $\mathbf{s}'$ is a Nash equilibrium. The cost of a job i is $2m - 1$. If that job is reassigned to any other machine, then the cost of i will be $2m$. Thus, $\mathbf{s}'$ is a Nash equilibrium. We will now calculate $c(\mathbf{s}')$. For each machine k, $x_k(\mathbf{s}') = m$ and $e_k(\mathbf{s}') = \frac{m(m-1)}{2}$. Thus, $c(\mathbf{s}') = 2m^3 - m^2$. Since $\frac{2m^3 - m^2}{m^3} = 2 - \frac{1}{m}$, we conclude that the price of anarchy is $\geq 2 - \frac{1}{m}$. ■

3.2 Strong Nash Equilibrium for $m = 2$

In this section we will examine strong Nash equilibrium for Balancing with Conflicts when $m = 2$. Note that games in which people partition themselves into two groups while trying to avoid being with others they dislike already present a nontrivial and problem [4, 14]. For our first result, we will show that Balancing with Conflicts always admits a strong Nash equilibrium when $m = 2$.

Theorem 3.7. The optimal solution for Balancing with Conflicts is a strong Nash equilibrium when $m = 2$.

Proof. Suppose, by way of contradiction, that the optimal solution $\mathbf{s}^*$ is not a strong Nash equilibrium. Then there is a deviating coalition in which every player in the coalition can decrease their cost. Call the solution that results from this deviation $\mathbf{s}'$. Let $X_{k,l} = \{i : s_i^* = k, s_i' = l\}$, $x_{k,l} = |X_{k,l}|$. Then the deviating coalition is the set $X_{1,2} \cup X_{2,1}$ and the set of players who do not deviate is $X_{1,1} \cup X_{2,2}$.

For each player $i \in X_{1,2}$, their cost decreases by $x_{1,1}+e(i, X_{1,1})$ and increases by $x_{2,2}+e(i, X_{2,2})$. Since i is in the deviating coalition, it must be the case that $x_{1,1}+e(i, X_{1,1}) > x_{2,2}+e(i, X_{2,2})$. Summing this up over all players $i \in X_{1,2}$ yields

$$x_{1,1}x_{1,2} + e(X_{1,1}, X_{1,2}) > x_{2,2}x_{1,2} + e(X_{2,2}, X_{1,2}).$$

Similarly, we can derive an inequality for each player $i \in X_{2,1}$. Summing this inequality over all players $i \in X_{2,1}$ results in

$$x_{2,2}x_{2,1} + e(X_{2,2}, X_{2,1}) > x_{1,1}x_{2,1} + e(X_{1,1}, X_{2,1}).$$

Next, we consider a player $i \in X_{1,1}$. Their cost decreases by $x_{1,2}+e(i, X_{1,2})$ and increases by $x_{2,1}+e(i, X_{2,1})$. Then $\sum_{i\in X_{1,1}} (c_i(\mathbf{s}') - c_i(\mathbf{s}^*)) = x_{1,1}x_{2,1} + e(X_{1,1}, X_{2,1}) - x_{1,1}x_{1,2} - e(X_{1,1}, X_{1,2})$.

Similarly, $\sum_{i\in X_{2,2}} (c_i(\mathbf{s}') - c_i(\mathbf{s}^*)) = x_{2,2}x_{1,2} + e(X_{2,2}, X_{1,2}) - x_{2,2}x_{2,1} - e(X_{2,2}, X_{2,1})$.

We can now calculate $c(\mathbf{s}') - c(\mathbf{s}^*)$.

$$\begin{aligned} c(\mathbf{s}') - c(\mathbf{s}^*) &= \sum_{i\in N} \big(c_i(\mathbf{s}') - c_i(\mathbf{s}^*)\big) \\ &= 2\,(x_{2,2}x_{1,2} + e(X_{2,2}, X_{1,2}) + x_{1,1}x_{2,1} + e(X_{1,1}, X_{2,1}) \\ &\quad -x_{1,1}x_{1,2} - e(X_{1,1}, X_{1,2}) - x_{2,2}x_{2,1} - e(X_{2,2}, X_{2,1})) \\ &< 0 \end{aligned}$$

by our inequalities, which is a contradiction. ∎

It immediately follows from this theorem that the strong price of stability is 1. We suspect that strong Nash equilibrium need not exist for $m \geq 3$, but we've yet to find any such examples. The existence of strong Nash equilibrium for cut games in general is unresolved, as seen in [14, 15].

Next, we will provide an upper bound for the strong price of anarchy with the same technique used to bound the strong price of anarchy for the max k-cut game in [15]. This technique fails when $m > 2$, partially because the technique in [15] is for max k-cut as opposed to min k-uncut, which is the type of cut problem that we are considering here.

Theorem 3.8. The strong price of anarchy is at most $\frac{4}{3} + \frac{2}{3n}$ Balancing with Conflicts with $m = 2, n \geq 4$. The strong price of anarchy is 1 for $m = 2, n \leq 3$.

Proof. Let $\mathbf{s}^*$ denote an optimal solution and let $\mathbf{s}'$ denote a permutation of $\mathbf{s}^*$, i.e., $X_1(\mathbf{s}^*) = X_2(\mathbf{s}')$ and $X_2(\mathbf{s}^*) = X_1(\mathbf{s}')$. Let $\mathbf{s}$ be a strong Nash equilibrium. Let $X_{k,l} = \{i : s_i = k, s_i^* = l\}$, $x_{k,l} = |X_{k,l}|$. Then

$$c(\mathbf{s}) = (x_{1,1}+x_{1,2})^2 + (x_{2,1}+x_{2,2})^2 + 2\left(\sum_{k=1}^{2}\sum_{l=1}^{2} e(X_{k,l}) + e(X_{1,1}, X_{1,2}) + e(X_{2,2}, X_{2,1})\right)$$

$$c(\mathbf{s}^*) = c(\mathbf{s}') = (x_{1,1}+x_{2,1})^2 + (x_{1,2}+x_{2,2})^2 + 2\left(\sum_{k=1}^{2}\sum_{l=1}^{2} e(X_{k,l}) + e(X_{1,1}, X_{2,1}) + e(X_{2,2}, X_{1,2})\right).$$

Let $X_{1,2} \cup X_{2,1}$ be a deviating coalition in $\mathbf{s}$. By the definition of $X_{k,l}$ and the fact there are only two strategies, the resulting state of this deviation is $\mathbf{s}^*$. However, since $\mathbf{s}$ is a strong Nash equilibrium, that must mean there is a player $i_1 \in X_{1,2} \cup X_{2,1}$ that does not want to deviate with the coalition. Furthermore, there must be a player i_2 that does not want deviate with the coalition $X_{1,2} \cup X_{2,1} \backslash \{i_1\}$. We can repeat this process for each player in $X_{1,2} \cup X_{2,1}$, which gives us an

inequality for each player. Suppose that player $i_t \in X_{1,2}$ is the t'th player to leave the coalition in this way. Let $Y^t_{1,2}$ denote the set of players in $X_{1,2}$ that left the coalition before i_t and $Y^t_{2,1}$ players in $X_{2,1}$ that left the coalition before i_t. Let $y^t_{1,2} = |Y^t_{1,2}|$ and $y^t_{2,1} = |Y^t_{2,1}|$. Then

$$c_{i_t}(\mathbf{s}) \le x_{2,2} + x_{1,2} - y^t_{1,2} + y^t_{2,1} + e\left(\{i_t\}, X_{2,2}\right) + e\left(\{i_t\}, X_{1,2} \backslash Y^t_{1,2}\right) + e\left(\{i_t\}, Y^t_{2,1}\right).$$

Similarly, if player i_t is in $X_{2,1}$, then

$$c_{i_t}(\mathbf{s}) \le x_{1,1} + x_{2,1} - y^t_{2,1} + y^t_{1,2} + e\left(\{i_t\}, X_{1,1}\right) + e\left(\{i_t\}, X_{2,1} \backslash Y^t_{2,1}\right) + e\left(\{i_t\}, Y^t_{1,2}\right).$$

Our goal is to find an upper bound for $\sum_{i \in X_{1,2} \cup X_{2,1}} c_i(\mathbf{s})$. We will do this by summing our previous inequalities over all players in $X_{1,2} \cup X_{2,1}$. We claim that doing this gives us

$$\begin{aligned} \sum_{i \in X_{1,2} \cup X_{2,1}} c_i(\mathbf{s}) \le{}& x_{2,2}x_{1,2} + x_{1,1}x_{2,1} + \frac{x^2_{1,2} + x_{1,2}}{2} + \frac{x^2_{2,1} + x_{2,1}}{2} + x_{1,2}x_{2,1} \\ &+ e(X_{2,2}, X_{1,2}) + e(X_{1,1}, X_{2,1}) + e(X_{1,2}) + e(X_{2,1}) + e(X_{1,2}, X_{2,1}). \qquad (1) \end{aligned}$$

To prove Inequality (1), we will consider each term separately. Clearly, $\sum_{i \in X_{1,2}} x_{2,2} = x_{2,2}x_{1,2}$ and $\sum_{i \in X_{2,1}} x_{1,1} = x_{1,1}x_{2,1}$. Similarly, we observe that $\sum_{i \in X_{1,2}} e(i, X_{2,2}) = e(X_{1,2}, X_{2,2})$ and $\sum_{i \in X_{2,1}} e(i, X_{1,1}) = e(X_{2,1}, X_{1,1})$.

Next, we consider $\sum_{i_t \in X_{1,2}} (x_{1,2} - y^t_{1,2})$. Consider the first player i_t in $X_{1,2}$ to leave the coalition. Then $y^t_{1,2} = 0$, since no player in $X_{1,2}$ has left the coalition before i_t. For the second player $i_{t'}$ in $X_{1,2}$ to leave the coalition, $y^{t'}_{1,2} = 1$, because player $i_t \in Y^{t'}_{1,2}$. If we continue to repeat this process, we observe that for the kth player i_t in $X_{1,2}$ to leave the coalition, $y^t_{1,2} = k - 1$. Thus, it follows that

$$\sum_{i_t \in X_{1,2}} \left(x_{1,2} - y^t_{1,2}\right) = \sum_{k=0}^{x_{1,2}-1} (x_{1,2} - k) = \frac{x^2_{1,2} + x_{1,2}}{2}.$$

Now we consider $\sum_{i_t \in X_{1,2}} e(\{i_t\}, X_{1,2} \backslash Y^t_{1,2})$. Consider an edge $(i_t, i_{t'})$ where $i_t, i_{t'} \in X_{1,2}$, $t < t'$. Since $i_{t'} \notin Y^t_{1,2}$, $(i_t, i_{t'}) \in E(i_t, X_{1,2} \backslash Y^t_{1,2})$ which means that the edge is counted for player i_t in the sum above. However, since $i_t \in Y^{t'}_{1,2}$, $(i_t, i_{t'}) \notin E(i_{t'}, X_{1,2} \backslash Y^{t'}_{1,2})$ which means that the edge is not counted for player $i_{t'}$ in this sum. Thus, we conclude that each edge in $E(X_{1,2})$ is counted exactly once. That is,

$$\sum_{i_t \in X_{1,2}} e\left(\{i_t\}, X_{1,2} \backslash Y^t_{1,2}\right) = e\left(X_{1,2}\right).$$

Similar analysis holds for players in $X_{2,1}$. That is, $\sum_{i_t \in X_{2,1}} \left(x_{2,1} - y^t_{2,1}\right) = \frac{x^2_{2,1} + x_{2,1}}{2}$ and $\sum_{i_t \in X_{2,1}} e\left(\{i_t\}, X_{2,1} \backslash Y^t_{2,1}\right) = e\left(X_{2,1}\right)$.

We will now calculate $\sum_{i_t \in X_{1,2}} y^t_{2,1} + \sum_{i_t \in X_{2,1}} y^t_{1,2}$. Consider a pair of players $i_t \in X_{1,2}$ and $i_{t'} \in X_{2,1}$. Suppose, without loss of generality, that i_t leaves the coalition first. Since $i_{t'} \notin Y^t_{2,1}$ and $i_t \in Y^{t'}_{1,2}$, this pair contributes 1 to our sum. Since this holds for every pair of players and there are $x_{1,2}x_{2,1}$ such pairs, we conclude that $\sum_{i_t \in X_{1,2}} y^t_{2,1} + \sum_{i_t \in X_{2,1}} y^t_{1,2} = x_{1,2}x_{2,1}$. Similarly, $\sum_{i_t \in X_{1,2}} e(\{i_t\}, Y^t_{2,1}) + \sum_{i_t \in X_{2,1}} e(\{i_t\}, Y^t_{1,2}) = e(X_{1,2}, X_{2,1})$.

Combining these results gives us (1).

We will now consider another deviating coalition, $X_{1,1} \cup X_{2,2}$, which results in the state $\mathbf{s}'$. The exact same analysis holds for this coalition, giving us

$$\sum_{i \in X_{1,1} \cup X_{2,2}} c_i(\mathbf{s}) \leq x_{2,2}x_{1,2} + x_{1,1}x_{2,1} + \frac{x_{1,1}^2 + x_{1,1}}{2} + \frac{x_{2,2}^2 + x_{2,2}}{2} + x_{1,1}x_{2,2} + e(X_{2,2}, X_{1,2}) + e(X_{1,1}, X_{2,1}) + e(X_{1,1}) + e(X_{2,2}) + e(X_{1,1}, X_{2,2}). \tag{2}$$

We can combine (1) and (2) to give us an upper bound on $c(\mathbf{s})$. Using this inequality and the value of $c(\mathbf{s}^*)$, we can derive

$$\begin{aligned}
c(\mathbf{s}) &\leq \sum_{k=1}^{2}\sum_{l=1}^{2}\left(\frac{x_{k,l}^2 + x_{k,l}}{2} + e\left(X_{u,v}\right)\right) + 2x_{2,2}x_{1,2} + 2x_{1,1}x_{2,1} + x_{1,2}x_{2,1} \\
&\quad + x_{1,1}x_{2,2} + 2e\left(X_{2,2}, X_{1,2}\right) + 2e\left(X_{1,1}, X_{2,1}\right) + e\left(X_{1,2}, X_{2,1}\right) + e\left(X_{1,1}, X_{2,2}\right) \\
&= \frac{1}{2}c(\mathbf{s}^*) + \frac{n}{2} + x_{2,2}x_{1,2} + x_{1,1}x_{2,1} + x_{1,2}x_{2,1} + x_{1,1}x_{2,2} \\
&\quad + e\left(X_{2,2}, X_{1,2}\right) + e\left(X_{1,1}, X_{2,1}\right) + e\left(X_{1,2}, X_{2,1}\right) + e\left(X_{1,1}, X_{2,2}\right) \\
&= \frac{1}{2}c(\mathbf{s}^*) + \frac{n}{2} + x_{2,2}x_{1,2} + x_{1,1}x_{2,1} + x_{1,2}x_{2,1} + x_{1,1}x_{2,2} \\
&\quad + |E| - e\left(X_{1,1}, X_{1,2}\right) - e\left(X_{2,1}, X_{2,2}\right) - \sum_{k=1}^{2}\sum_{l=1}^{2} e\left(X_{k,l}\right)
\end{aligned}$$

Adding $\frac{1}{2}c(\mathbf{s})$ to both sides, and applying Lemmas 3.4 and 3.5 allows us to derive

$$\begin{aligned}
\frac{3}{2}c(\mathbf{s}) &\leq \frac{1}{2}c(\mathbf{s}^*) + \frac{n}{2} + \frac{1}{2}\sum_{k=1}^{2}\sum_{l=1}^{2} x_{k,l}^2 + x_{1,1}x_{1,2} + x_{2,1}x_{2,2} \\
&\quad + x_{2,2}x_{1,2} + x_{1,1}x_{2,1} + x_{1,2}x_{2,1} + x_{1,1}x_{2,2} + |E| \\
&= \frac{1}{2}c(\mathbf{s}^*) + \frac{n}{2} + \frac{1}{2}(x_{1,1} + x_{1,2} + x_{2,1} + x_{2,2})^2 + |E| \\
&= \frac{1}{2}c(\mathbf{s}^*) + \frac{n}{2} + \frac{n^2}{2} + |E| \\
&\leq \left(2 + \frac{1}{n}\right) c(\mathbf{s}^*)
\end{aligned}$$

which completes the proof. ■

Unfortunately, this bound is not tight; the best lower bound we have found for the strong price of anarchy is $\frac{5}{4}$, using the same example as the one used in [14].

Claim 3.9. The strong price of anarchy of Balancing with Conflicts with $m = 2$ is at least $\frac{5}{4}$.

Proof. Let G be a path on four nodes. The solution in which two jobs are assigned to each machine such that no job incurs edge cost is the optimal solution with $c(\mathbf{s}^*) = 8$.

Consider the solution $\mathbf{s}$ formed by assigning the two degree one jobs to one machine and the other two jobs to the other machine. Then the degree one jobs each have cost 2 and the other two jobs have cost 3, which means $c(\mathbf{s}) = 10$. Since neither of the jobs with cost 2 will join a deviating coalition, this solution is a strong Nash equilibrium. ■

3.3 Extensions

3.3.1 Weighted Edges

Naturally, we would want to generalize our model to include edges with arbitrary weights. However, because the price of anarchy of the min uncut game is unbounded, adding edge weights to Balancing with Conflicts results in the price of anarchy approaching infinity. Consider an extension of Balancing with Conflicts in which the maximum weight of an edge is w. Then using the exact same arguments as our original model, we can derive the following bounds for the price of stability and price of anarchy.

Theorem 3.10. The price of stability for Balancing with Conflicts with maximum edge weight w is 1.

Theorem 3.11. The price of anarchy for Balancing with Conflicts with maximum edge weight w is $1 + w - \frac{w}{m} + \frac{m-1}{n}$.

3.3.2 Balancing with Friendship

Definition. We now consider a variation of Balancing with Conflicts in which the additional cost arises from separating players who are "friends". The Balancing with Friendship game is identical to the Balancing with Conflicts except jobs incur additional costs for having edges with jobs that are not assigned to the same machine. That is, for every job i with $s_i = k$, $c_i(\mathbf{s}) = x_k(\mathbf{s}) + \sum_{l \neq k} e(\{i\}, X_l(\mathbf{s}))$.

We can use the exact same techniques for the Balancing with Friendship game as we did with the Balancing with Conflicts game to prove better results: we are able to obtain tight price of anarchy bounds for Balancing with Friendship. As with the Balancing with Conflicts game, the Balancing with Friendship game is an exact potential game with $\Phi(\mathbf{s}) = \frac{1}{2}c(\mathbf{s})$, which implies that the price of stability is 1.

Theorem 3.12. The Balancing with Friendship game is $\left(2 - \frac{1}{m}, 0\right)$-semi-smooth. The price of total anarchy is at most $2 - \frac{1}{m}$, and it is tight.

Proof. We can use the same argument as above to show that Balancing with Friendship is semi-smooth if we use the following two lower bounds for OPT: for any solution $\mathbf{s}$, $c(\mathbf{s}) \geq \frac{n^2}{m}$, and $c(\mathbf{s}) \geq 2|E| + n$.

To show our bound is tight, let G consist of m disjoint cliques of size m each. Assigning each clique to the same machine such that only one clique is assigned to each machine is the optimal solution. Assigning each job in the same clique to a different machine is the worst case Nash equilibrium. ■

Corollary 3.13. For any instance of the Balancing with Friendship and for any $\epsilon > 0$, the BR dynamic reaches a state $\mathbf{s}^t$ with

$$c(\mathbf{s}^t) \leq \left(2 - \frac{1}{m}\right)(1 + \epsilon) \cdot c(\mathbf{s}^*)$$

in at most $O\left(n \log \frac{m}{\epsilon}\right)$ steps from any initial state $\mathbf{s}^0$. This approximation factor holds for all further states reached by best-response dynamics.

Theorem 3.14. The optimal solution is a strong Nash equilibrium for Balancing with Friendship with $m = 2$, and the strong price of anarchy is at most $\frac{4}{3}$.

3.3.3 Balancing with Conflicts and Friendship

Balancing with Conflicts and Friendship is essentially a combination of Balancing with Conflicts and Balancing with Friendship; the graph G now has two types of edges, friendship edges and enemy edges. Jobs incur additional costs for sharing enemy edges with jobs that are assigned to the same machine and for sharing friendship edges with jobs that are assigned to different machines. Using the same techniques, we can show that Balancing with Conflicts and Friendship is $\left(3 - \frac{2}{m}, 0\right)$-semi-smooth.

4 Sharing with Conflicts

Definition. We are given a set of players $N = \{1, 2, \ldots, n\}$, a set of machines $M = \{1, 2, \ldots, m\}$ with positive values $P = \{p_1, p_2, \ldots, p_m\}$, and $G = (N, E)$, an undirected graph on the player set. Each player must assign a job to a single machine to be processed.

We will use the same notation as we did in Balancing with Conflicts. Each machine k will distribute the p_k utility associated with it equally among the jobs assigned to it. Each job will receive 1 utility for each edge e that it shares with another job that is assigned to another machine. We define the utility of job i with $s_i = k$ as $u_i(\mathbf{s}) = \frac{p_k}{x_k(\mathbf{s})} + \sum_{l \neq k} e(\{i\}, X_l(\mathbf{s}))$. Finally, we define the utility of a solution $\mathbf{s}$, denoted $u(\mathbf{s})$, as the sum of the utility of all jobs. That is, $u(\mathbf{s}) = \sum_{i \in N} u_i(\mathbf{s})$.

4.1 Tight Bounds for Quality of CCE

We will now provide analysis of the Sharing with Conflicts game. Observe that Sharing with Conflicts is an exact potential game with $\Phi(\mathbf{s}) = \sum_{k=1}^{m} \sum_{l=1}^{x_k(\mathbf{s})} \frac{p_k}{l} + \sum_{k=1}^{m-1} \sum_{l=k+1}^{m} e\left(X_k(\mathbf{s}), X_l(\mathbf{s})\right)$, which guarantees the existence of pure Nash equilibria. We will now prove the main result of this section: Sharing with Conflicts is semi-smooth. This, along with our lower bound example, will give tight price of anarchy bounds for many equilibrium concepts.

Theorem 4.1. Sharing with Conflicts is $\left(\frac{m-1}{m}, \frac{m-2}{m}\right)$*-semi-smooth.*

Proof. Let $\mathbf{s}$ be an outcome, and let d_i denote the degree of player i in G. We will assume that $n \geq m$, but this argument works similarly for $n < m$. Suppose, without loss of generality, that only the first t machines have players assigned to them in $\mathbf{s}$. Then $u(\mathbf{s}) \geq \sum_{k=1}^{t} p_k$. For each player i with $s_i = k$, $u_i(\mathbf{s}) = p_k \cdot x_k(\mathbf{s})^{-1} + \sum_{l \neq k} e(\{i\}, X_l(\mathbf{s}))$. If player i deviates to a different machine l, then their utility will be the sum of p_l divided by the number of players already on machine l plus 1 for itself and the number of edges from i to players not on l. That is, for each $s'_i = l \neq s_i$, $u_i(s'_i, s_{-i}) = p_l \cdot (x_l(\mathbf{s}) + 1)^{-1} + \sum_{k \neq l} e(\{i\}, X_k(\mathbf{s}))$. For each player i, let σ_i denote the mixed strategy in which every strategy $s'_i \in S_i$ is selected with probability $\frac{1}{m}$. Then using the fact that $x_l(\mathbf{s}) = 0$ for $l > t$ allows us to derive that for every player i with $s_i = k$,

$$\mathbf{E}[u_i(\sigma_i, s_{-i})] = \frac{1}{m} \left(\sum_{l=1 : l \neq k}^{m} \frac{p_l}{x_l(\mathbf{s}) + 1} + \frac{p_k}{x_k(\mathbf{s})} + \sum_{l=1}^{m} \sum_{l' \neq l} e(\{i\}, X_{l'}(\mathbf{s})) \right).$$

It is easy to see that $\sum_{l=1}^{m} \sum_{l' \neq l} e(\{i\}, X_{l'}(\mathbf{s})) = (m-1)d_i$, since all edges from i to each machine l appear $m-1$ times in the sum. Notice also that for machines $l > t$, we know that $x_l(s) = 0$ by

definition of t, and so $\frac{p_l}{x_l(\mathbf{s})+1} = p_l$. Thus,

$$\begin{aligned}
\mathbf{E}[u_i(\sigma_i, s_{-i})] &= \frac{1}{m}\left(\sum_{l=1:l\neq k}^{t} \frac{p_l}{x_l(\mathbf{s})+1} + \sum_{l=t+1}^{m} p_l + \frac{p_k}{x_k(\mathbf{s})} + (m-1)d_i\right) \\
&\geq \frac{1}{m}\left(\frac{p_k}{x_k(\mathbf{s})} + \sum_{l=t+1}^{m} p_l + (m-1)d_i\right).
\end{aligned}$$

By linearity of expectation and the fact that $\sum_{i:s_i=k} \frac{p_k}{x_k(\mathbf{s})} = p_k$,

$$\begin{aligned}
\mathbf{E}\left[\sum_{i\in N} u_i(\sigma_i, s_{-i})\right] &= \sum_{i\in N} \mathbf{E}[u_i(\sigma_i, s_{-i})] \\
&\geq \frac{1}{m}\left(\sum_{l=1}^{t} p_l + n\sum_{l=t+1}^{m} p_l + 2(m-1)|E|\right)
\end{aligned}$$

It is clear that for any solution $\mathbf{s}'$, $u(\mathbf{s}') \leq \sum_{l=1}^{m} p_l + 2|E|$. Using this, our lower bound for $u(\mathbf{s})$, and the fact that $n \geq m$, we can show our previous quantity is at least

$$\begin{aligned}
\frac{1}{m}\left(\sum_{l=1}^{t} p_l + 2(m-1)|E|\right) + \frac{m-1}{m}\sum_{l=t+1}^{m} p_l &\geq \frac{m-1}{m}\left(\sum_{l=1}^{m} p_l + 2|E|\right) - \frac{m-2}{m}\sum_{l=1}^{t} p_l \\
&\geq \frac{m-1}{m}u(\mathbf{s}^*) - \frac{m-2}{m}u(\mathbf{s}).
\end{aligned}$$

The same argument works when $n < m$ because we can ignore the $m-n$ machines with the least value. Players only need to deviate to the n machines with the highest value, and the upper bound we obtain for OPT is lower since at most n machines can be covered. ■

We will now give a lower bound for the price of stability, which also acts as a lower bound for the price of anarchy, to show that the upper bound we found with semi-smoothness is tight.

Claim 4.2. The price of stability of the Sharing with Conflicts game is at most 2, and it is tight.

Proof. Let $n = m$, $G = K_m$, $p_1 = m^2 - m + \epsilon$ for any $\epsilon > 0$ and for all machines $k \neq 1$, $p_k = 0$. The optimal solution is to assign each job to a different machine, in which case every job receives $m-1$ utility from edges and the job assigned to machine 1 receives an additional $m^2 - m + \epsilon$ utility from the machine. Thus, $u(\mathbf{s}^*) = 2m^2 - 2m + \epsilon$.

Any player is guaranteed to receive at least $m-1+\frac{\epsilon}{m}$ utility if they choose machine 1. However, each player can receive at most $m-1$ utility if they choose any of the other machines. We conclude that the solution $\mathbf{s}$ formed by assigning all jobs to machine 1 is the only pure Nash equilibrium. We observe that $u(\mathbf{s}) = m^2 - m + \epsilon$ which approaches $\frac{1}{2}u(\mathbf{s}^*)$ as $\epsilon \to 0$.

Theorem 4.1 guarantees that the price of anarchy is at most 2, which implies the price of stability cannot exceed 2. ■

This same example also shows that the strong price of stability is 2. The following price of anarchy results follow immediately from Theorem 4.1 and our lower bound example from Claim 4.2:

Corollary 4.3. The price of total anarchy is at most 2, and it is tight. The same holds for the strong price of anarchy.

This result is interesting because the price of anarchy for the market sharing game with m machines is $2 - \frac{1}{m}$ and the price of anarchy of the max k-cut game is $\frac{k}{k-1}$. For $m > 2$, this is lower than the price of stability (and anarchy) of Sharing with Conflicts. Since our game is essentially a combination of these two games, this means that combining two games can lead to greater inefficiency.

Finally, we will consider convergence for the Sharing with Conflicts game. Our convergence results for Sharing with Conflicts are weaker than those for Balancing with Conflicts because Sharing with Conflicts is not a perfect game, which means we cannot use the convergence result from [1]. Instead, we use a generalization of the convergence result from [27] to obtain our first result: that a good solution is reached very quickly. Then we use a modified version of the convergence result from [1] that does not require the game to be perfect. The trade-off is that the approximation factor of the solutions reached is not as good. We observe that $\log(n) \cdot \Phi(\mathbf{s}) \geq u(\mathbf{s}) \geq \frac{1}{2}\Phi(\mathbf{s})$ for this game. Thus, by Theorems 4.1, 2.6, and 2.7, we have the following two convergence results.

Corollary 4.4. For any Sharing with Conflicts instance and for any $\epsilon > 0$, the BR dynamic reaches a state $\mathbf{s}^t$ with $u(\mathbf{s}^t) \geq \frac{1}{2+2\epsilon} \cdot u(\mathbf{s}^*)$ from any starting state $\mathbf{s}^0$ in at most $O\left(\frac{n}{\epsilon} \log u(\mathbf{s}^*)\right)$ steps.

Corollary 4.5. For any Sharing with Conflicts instance and for any $\epsilon > 0$, the BR dynamic converges to a state $\mathbf{s}^t$ with $u(\mathbf{s}^t) \geq \frac{1-\epsilon}{4\log(n)} \cdot u(\mathbf{s}^*)$ in at most $O\left(\frac{n}{\log(n)} \log \frac{1}{\epsilon}\right)$ steps from any initial state. Furthermore, all future states reached with best-response dynamics will satisfy this approximation factor as well.

4.2 Extensions

4.2.1 Weighted Edges

All of our results in the previous section hold using the same arguments if we allow the edges to have positive weights and redefine the utility function of a job i such that it receives w_e utility for each edge e that it shares with another job that is assigned to another machine. Thus, unlike in Section 3, the bounds on the efficiency of equilibrium remain small (at most 2) even in the presence of weighted edges.

4.2.2 Sharing with Friendship

Definition. Just as we considered a utility version of the Balancing with Conflicts game, we will now consider a utility version of the Balancing with Friendship game. The Sharing with Friendship game is the same as the Sharing with Conflicts game except that jobs receive additional utility for sharing edges with jobs on the same machine. That is, for every job i with $s_i = k$, $u_i(\mathbf{s}) = \frac{p_k}{x_k(\mathbf{s})} + e(\{i\}, X_k(\mathbf{s}))$.

We will prove similar results for the Sharing with Friendship using the same arguments as we did for the Sharing with Conflicts. Furthermore, all of these results hold if we allow arbitrary positive edge weights. Sharing with Friendship is an exact potential game with $\Phi(\mathbf{s}) = \sum_{k=1}^{m} \sum_{l=1}^{x_k(\mathbf{s})} \frac{p_k}{l} + \sum_{k=1}^{m} e_k(\mathbf{s})$.

Theorem 4.6. The Sharing with Friendship is $\left(\frac{1}{m}, 0\right)$*-semi-smooth.* The price of total anarchy is at most m, and it is tight. The price of stability is at least $2 - \frac{1}{m}$.

Corollary 4.7. For any Sharing with Friendship instance and for any $\epsilon > 0$, the BR dynamic reaches a state $\mathbf{s}^t$ with $u(\mathbf{s}^t) \geq \frac{1}{m(1+\epsilon)} \cdot u(\mathbf{s}^*)$ from any starting state $\mathbf{s}^0$ in at most $O\left(\frac{n}{\epsilon} \log u(\mathbf{s}^*)\right)$ steps.

Corollary 4.8. For any Sharing with Friendship instance and for any $\epsilon > 0$, the BR dynamic reaches to a state $\mathbf{s}^t$ with $u(\mathbf{s}^t) \geq \frac{(1-\epsilon)}{2m\log(n)} \cdot u(\mathbf{s}^*)$ in at most $O\left(\frac{n}{\log(n)} \log \frac{1}{\epsilon}\right)$ steps from any initial state. Furthermore, all future states reached with best-response dynamics will satisfy this approximation factor as well.

Unlike in Sharing with Conflicts, where existence of strong Nash equilibrium is an open question even for $m = 2$, we can show that strong Nash equilibrium does not necessarily exist for this game. The strong price of anarchy is known to be at most 3 [3], and our price of stability example gives a lower bound of 2.

Theorem 4.9. Strong Nash equilibrium need not exist for the Sharing with Friendship game.

Proof. Consider an instance in which G is a graph with 4 nodes and the edges $(1, 2)$ and $(3, 4)$. There are two machines with values $2 + \epsilon$ and $4 + 3\epsilon$ for some $\epsilon > 0$. It is easily verified that no strong Nash equilibrium exists for this instance. ■

5 Conclusion and Future Work

In this paper, we introduce two new job-assignment games – Balancing with Conflicts, and Sharing with Conflicts – which allow us to model positive and negative interactions between certain pairs of jobs. For both games, we provide tight or nearly-tight bounds on the Price of Anarchy for pure Nash, mixed Nash, correlated, coarse correlated, and strong Nash equilibria. We also bound the time needed for convergence to near-optimal outcomes. For most of these results, we make use of the notion of semi-smoothness [7] and observe that standard smoothness [27] is not sufficient for achieving tight bounds.

There are a number of natural avenues for further study. Some of the more direct extensions include generalizing our existing results to broader classes of latency functions; in the case of balancing with conflicts, proving tight bounds on the price of anarchy seems difficult even for related machines, and may require new techniques.

Another promising direction is to improve bounds on the quality of equilibria when the conflict graph's structure is constrained. For example, consider our motivating scenario in which jobs conflict over the use of certain computational resources. If the number of distinct resource types is small, the set of possible conflict graphs is limited and thus better bounds may apply. More generally, it would be interesting to find any structural parametrization that increases equilibrium quality.

Perhaps most interesting would be a broader exploration of the process of combining games. All of the games studied in this paper can be viewed as a "sum" of simpler games. Given any pair of games $\mathcal{G}_1$ and $\mathcal{G}_2$ in which the players and their strategy sets are the same, we can define a new game $\mathcal{G}_1 \oplus \mathcal{G}_2$ with the same players, the same strategies, and utilities given by $u_i(\mathbf{s}) = u_i^1(\mathbf{s}) + u_i^2(\mathbf{s})$, where u_i^j is i's utility function in $\mathcal{G}_j$. In particular, the games we introduced are sums of load balancing games, market sharing games, and cut games. To what extent do these new games inherit properties of their component games? Some relationships are straightforward – if $\mathcal{G}_1$ and $\mathcal{G}_2$ are potential games, then clearly so too is $\mathcal{G}_1 \oplus \mathcal{G}_2$. Similarly, one can bound the smoothness of $\mathcal{G}_1 \oplus \mathcal{G}_2$ in terms of that of $\mathcal{G}_1$ and $\mathcal{G}_2$ (although the naive bounds do not appear to be tight). Under what conditions and to what extent are various game characteristics (such as smoothness) preserved under this game operator? Are there natural classes of games that are well behaved under this operation? The same questions can be applied to other game operators as well, such as product, min, and max. The development of tools for analyzing such operators is a promising direction for future research.

Acknowledgements

We thank Carlos Varela for illuminating discussions on the subject of job assignment.

References

[1] J. Augustine, N. Chen, E. Elkind, A. Fanelli, N. Gravin, and D. Shiryaev. Dynamics of profit-sharing games. In *IJCAI '11*, pages 37–42.

[2] B. Awerbuch, Y. Azar, A. Epstein, V. S. Mirrkoni, and A. Skopalik. Fast convergence to nearly optimal solutions in potential games. In *EC '08*, pages 264–273.

[3] Y. Bachrach, V. Syrgkanis, and M. Vojnović. Efficiency and the redistribution of welfare. Working Paper, 2011.

[4] A. Bhalgat, T. Chakraborty, and S. Khanna. Approximating pure nash equilibrium in cut, party affiliation, and satisfiability games. In *EC '10*, pages 73–82.

[5] A. Blum, M. Hajiaghayi, K. Ligett, and A. Roth. Regret minimization and the price of total anarchy. In *STOC '08*, pages 373–382.

[6] I. Caragiannis, M. Flammini, C. Kaklamansis, P. Kanellopoulos, and L. Moscardelli. Tight bounds for selfish and greedy load balancing. In *ICALP '06*, pages 311–322.

[7] I. Caragiannis, C. Kaklamansis, P. Kanellopoulos, M. Kyropoulou, B. Lucier, R. P. Leme, and É. Tardos. On the efficiency of equilibria in generalized second price auctions. arXiv:1201.6429

[8] G. Christodoulou and E. Koutsoupias. The price of anarchy of finite congestion games. In *STOC '05*, pages 67–73.

[9] G. Christodoulou, V. S. Mirrokni, and A. Sidiropoulos. Convergence and approximation in potential games. In *STACS '06*, pages 349–360.

[10] A. Fanelli, M. Flammini, and L. Moscardelli. The speed of convergence in congestion games under best-response dynamics. In *ICALP '08*, pages 796–807.

[11] A. Fanelli and L. Moscardelli. On best response dynamics in weighted congestion games with polynomial delays. In *WINE '09*, pages 55–66.

[12] M. Feldman, L. Lewin-Eytan, and J. S. Naor. Hedonic clustering games. In *SPAA '12*, pages 267–276.

[13] M. X. Goemans, L. Li, V. S. Mirrokni. and M. Thottan. Market sharing games applied to content distribution in ad-hoc networks. In *JSAC '06*, 24(5):1020–1033.

[14] L. Gourvès and J. Monnot. On strong equilibria in the max cut game. In *WINE '09*, pages 608–615.

[15] L. Gourvès and J. Monnot. The max k-cut game and its strong equilibria. In *TAMC '10*, pages 234–246.

[16] M. Hoefer. Cost Sharing and Clustering under Distributed Competition. PhD Thesis, Universität Konstanz, 2007.

[17] M. O. Jackson and L. Yariv. Diffusion of behavior and equilibrium properties in network games. *American Economic Review*, 97(2):92–98, 2007.

[18] D. Kempe, J. Kleinberg, and É. Tardos. Maximizing the spread of influence through a social network. In *KDD '03*, pages 137–146.

[19] J. Kleinberg. Cascading behavior in networks: algorithmic and economic issues. *Algorithmic Game Theory*, N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors, pages 613–632, 2007.

[20] A. Kothari, S. Suri, C. D. Tóth, and Y. Zhou. Congestion games, load balancing, and the price of anarchy. In *CAAN '04*, pages 13–27.

[21] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *STACS '99*, pages 404–413.

[22] B. Lucier and R. P. Leme. GSP auctions with correlated types. In *EC '11*, pages 71–80.

[23] V. S. Mirrokni and A. Vetta. Convergence issues in competitive games. In *RANDOM-APPROX '04*, pages 183–194.

[24] D. Monderer and L. S. Shapley. Potential games. *Games and Economic Behavior*, 14(1):124–143, 1996.

[25] S. Morris. Contagion. *The Review of Economic Studies*, 67(1):57–78, 2000.

[26] N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007.

[27] T. Roughgarden. Intrinsic robustness of the price of anarchy. In *STOC '09*, pages 513–522.

[28] T. Roughgarden and F. Schoppmann. Local smoothness and the price of anarchy in atomic splittable congestion games. In *SODA '11*, pages 255–267.

[29] S. Suri, C. D. Tóth, and Y. Zhou. Selfish load balancing and atomic congestion games. *Algorithmica*, 47(1):79–96, 2007.

[30] B. Vöcking. Selfish load balancing. *Algorithmic Game Theory*, N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors, pages 517–542, 2007.